Oral History Panel on IBM 3380 Disk Drive

Interviewed by:
Jim Porter

Recorded: January 3, 2006
Mountain View, California

CHM Reference number: X3413.2006

Abstract: *The IBM 3380, announced in June, 1980, and eventually shipped in October, 1981, for mainframe applications, encountered many difficult development problems, but eventually became one of IBM's all-time highest cash producers. Panel members played a key role in solving difficult problems with recording heads, disks, motors, etc., and in preparing the drive models with increased capacity which were introduced through the 1980s.*

Jim Porter: We're here to discuss one of the most interesting disk drive development and production programs in the history of the computer industry and certainly of IBM, the IBM 3380, which was a disk drive used back in the good old mainframe days, which by the way, are still here to a certain extent as we sit here in 2006. But the 3380, which was announced in June 11, 1980, wasn't actually shipped until October of 1981 were the first shipments, representing in that delay of time, the difficulty in getting this job done. It was a major project, it was a very advanced disk drive for its day. And it became eventually one of the biggest cash producers for IBM, I think, in their history in that time. It was a drive which caused great difficulty for IBM's competitors. And for IBM's users it provided a much more cost effective storage set of tools than they had had up until that point in time. So to get started today, we have five of the people who made all of that possible at IBM, and let's start by asking each of these gentlemen to give us a little background on their own personal history.

Dick Whitney: Thank you, Jim. I'm Dick Whitney. I started with IBM in 1962 in the field as assistant engineer, and later as a salesman. Came to San Jose in 1969 to work on the planning, both product and business planning for large disk storage systems. The first project that I worked on was the 3330, which was a large, for the time, a large capacity disk storage subsystem. After that we had a double-capacity version called the 3330 Model 11, 200 megabytes per disk drive. And after that I worked on the 3350, which was, again, a large capacity disk drive system, derived from the Winchester technology. And that product, the 3350 was the predominant large capacity disk storage used on large systems up to the time of the 3380. And so my background preceding the 3380 was all to do with storage for large mainframe systems. Chris?

Chris Coolures: And I'm Chris Coolures, and I joined IBM in 1960, and my first program I worked on was the 1301. And then worked on the Winchester file, started the Madrid program, the 3350. And then worked on the 3800 printer, and in '78 I was made the program manager of the 3380 system, the drive portion. And I was program manager from '78 through 1981 when it finally-- we were able to ship the product.

Jack Grogan: I'm Jack Grogan. I started with IBM in 1959, and retired in 2002. My product experience is over the 305 RAMAC, and the 1405, 1301, 3330, 3350, 3380, 3390, and then on to the SCSI-based products. During that time, my focus primarily was in channel integration. That is, the design of the recording channel, selection of heads, disks and that sort of thing. And I was the recording channel manager for 3380 for about four years.

Jim Lucke: Thanks, Jack. I'm Jim Lucke, and I joined IBM in 1970. I came straight out of school from UC Berkeley, joined the Winchester program back with Chris and Jack. And then worked on the Winchester technology, the 3350, and then back in 1977, we started switching people among different programs, so I moved over to the 3380 program. Then in 1982, I was the program manager of the 3380. In '84, I was the lab director. And since that time, I've held a variety of other jobs, plant manager, launch manager, involved in introduction of the products in Brazil and Vimecarti. And I'm one of the few people that're still left at IBM. So I'm still working there. So, Mike?

Mike Warner: I'm Mike Warner. I joined IBM in 1965. And where all these guys either are still there or retired from IBM, I left IBM in 1985. I worked on a whole series of products at IBM, 2311, 2314, 3350, 3380, so forth, prior to leaving IBM. I was mostly involved with heads and media, although at one time for a while I was the product manager for the 3380, and other high-end drives. After I left IBM, I worked for Maxstor. Micropolis as VP of engineering. I had a company of my own for a short time called ORCA technology. And for the last 10 or 11 years I've been working at a company called Tessera doing semi-conductor packaging in miniaturizing electronics. Jim?

Porter: Mike, since you're sitting down at the end, why don't you do just a very brief description, point out what's going on with that head disk assembly for the 3380 sitting there on the table.

Warner: Did anybody mention what the capacity of this thing is yet?

Porter: Well, it started out originally as 1.2 gigabytes.

Warner: Now that's 1.2 gigabytes was for a machine that actually had two of these units in.

<overlapping conversation>

Porter: 1.2 gigabytes was the capacity of each HDA.

Warner: And this is an HDA head disk assembly that has-- it's difficult to see, but there are nine disks in here. There's two actuators on either side. So one actuator is accessing, what we call the front half, the other actuator is accessing the back half. On each of these actuators are four arms, and each arm has two magnetic heads. So one head is going through the outer portion of the disk, and the other head is traversing the inner portion of the disk. So there're 15 data heads, and one servo head in this actuator assembly. And I'm sure we'll get into it in this discussion, but there're an enormous number of technical undertakings to get these to work. So at the time it was the biggest, fastest, most expensive disk drive. I think the box sold for like 100...

Porter: Well, the box, which consisted of two...

Warner: Two.

Porter: Two of those HDAs had a sales price of $101,550.

Warner: Now the A box was more, right? Was it the A box?

Porter: Yes.

Warner: About $120,000. And for these, and they stood in a meter wide, a meter deep, and two meter high assembly.

Grogan: Called the refrigerator size.

Warner: Yeah, about a big refrigerator, a big, deep refrigerator. So it was extremely difficult to make this large a device, with all its mechanical complications reach the aerial densities and the technical objectives that we had. And we'll go through that in the next hour or so.

Porter: Okay, thank you. We wanted to have Whitney, and Coolures discuss the market, the business perspectives and so forth in dealing with this kind of product. Why was it designed this way?

<laughter>

Coolures: Well, one, from the design of, let's say, the HDA it was really to get an actuator that had a balanced symmetrical design, and that created some real difficulties in the carriage, and the system was different. So from a device point of view it was in one way more complex mechanically, but from a dynamics point of view, it did have a lot of advantages. Now from the configuration itself, we were looking to have a device that was migratable, so you could migrate your data from a 3350 to 3380. So it had to be compatible that way, and that's why we had 15 tracks that were similar. The string configuration, we offered new attachment through the controller and a new control storage unit, which provided dual path, which we called dynamic pathing, I think, at that time, so that even though there's two actuators on a HDA, those were addressable, independently through the different paths. So that gave you good accessibility. I don't know, Dick, if you want to elaborate from a subsystem point of view.

Whitney: Well, in some ways the configuration of the product was driven more by the technology itself than by the marketplace. On the other hand, the capacity and the features of the machine were truly driven by the marketplace. Back at that time in the late ’70s, as I mentioned earlier, the 3350 had become the dominant large system storage device. And it had been a departure from the predecessor, in that the media was not removable. So the large systems customers had become accustomed to fixed storage, if you will, and the market was growing substantially. I think it’s important to remember that our principal customer were large mainframes. Most of those were IBM or IBM-compatible mainframes. So the people in Poughkeepsie who designed and built those mainframes depended on us to provide them storage that would make those systems competitive. Now, not only as you know, Jim, not only was the disk storage very competitive, but the systems area was very competitive. At that time, as I remember, Amdahl was producing compatible IBM systems, then National Semiconductor formed a company called NAS, I believe, National Advanced Systems, which merged or acquired Amdahl, which ultimately became Hitachi Data Systems, I believe. And so the large systems marketplace was very competitive. And the people in IBM who were producing large systems wanted the direct access storage to be a competitive feature for the system.

Porter: And while you’re mentioning the competition, there were some disk drive competitors in that 3350 era that were pretty aggressive also.

Whitney: Yes, and they were becoming more and more aggressive.

Porter: Storage Technology, and Memorex and so forth, were making 3350-equivalent drives. And becoming really tough competitors for IBM, weren’t they?

Whitney: That’s right. And so we were confronted with declining opportunity for the 3350, and wanted to have a very competitive product in the replacement, which the 3380 became. Now there were a couple of, I guess, overriding issues. Always the principle competitive major for disk storage was dollars per megabyte, but what was also happening was the computer rooms were filling up so that megabytes per square foot became a measure of significance and great importance to large systems users. So that-- back to your original question, what drove the configuration. It was important to have a lot of capacity in the smallest possible footprint. And to what degree that influenced the design, I can’t be sure. But it was one of the factors that we really thought about.

Grogan: Well, I think there was some pressure.

Coolures: You know, there was the concern about power, too, that...

Grogan: There was some pressure developed by Storage Tech announcing a double density 3350.

Whitney: That's right.

Grogan: Is that right?

Whitney: Yes.

Grogan: And I remember, they advertised megabytes per BTU, because cooling was getting to be a serious problem.

Warner: Well, I think it actually had a big difference on the 3380 configuration. The predecessor to the 3380 was -- I don't even remember the code names, but it started off as an 18-inch filler disks with actuators positioned around it, which would now take a much large space...

Coolures: Much bigger footprint.

Warner: It was a horizontally oriented disk drive.

Porter: But that was never introduced.

Warner: It was never introduced, but it exercised the thin film head technology, and a thin film disk technology, which then were adopted into the HDA structure, you know, that was vertically integrated. And two, in one box to help save space.

Lucke: Yeah, the other thing that was a focus was previous actuators were two together on one side. This one had coming from both sides, and that did some decoupling, which helped with some of the vibration problems we had.

Warner: Yep.

Porter: Well, looking at it from the point of view of the customer, as several of you talked about, during that era, as a publisher of the *Disk/Trend Reports*, which tracked all the industry, and previous to that, when I worked for one of those companies called Memorex. When I went to work for Memorex in '68, the first thing I did was I went out and traveled at least one day each with every one of the 45 salesmen they had in the United States, from Midtown Manhattan, where the salesmen who didn't even have a car, to drive all day territories in Texas. I was in hundreds and hundreds of glass houses, as we used to call them, computer rooms, raised floor, carefully isolated, closed off against smoking and all those contamination things. And I have to tell you that you couldn't be more accurate in talking about the value of the square footage. Because as they filled up the computer rooms with disk drives and other devices, the cost of expanding them became huge for most of those companies. So when IBM said, "Hey, we can give you three times the amount of storage or more, depending upon the configuration, per square foot of computer room...," boy! What a lot of smiles those salesmen got on those calls, huh?

Whitney: You highlight a point about the difference in the configuration of the 3350, the predecessor product, from the 3380. 3350 was tabletop high. Whereas the 3380 then became...

Coolures: It was double.

Whitney: It was twice that.

Warner: Yeah, two meters high.

Whitney: And consequently we got much...

Coolures: It was four times the capacity.

Whitney: Much better capacity per square foot.

Coolures: Same basic footprint or less.

Whitney: Interesting sidelight, at the same time, some of you probably remember, all the computer rooms, as Jim suggested, had raised floors, typically eight-inch or 12-inch raised floors, and those things became so congested with cables, it was virtually becoming solid copper under those raised floors. And it was very hard to move things around, to get things in, and so the simplicity that was represented in this particular design, I think, was very attractive.

Porter: So the customer was very happy with the offering, I think, most of the time. But what did it take to get this thing built? We were going to talk about the actual configuration of the machine. Mike showed a little bit of the outline with that HDA, but the transition of the kinds of heads, the kinds of media, the kinds of channels, etcetera. What were the biggest challenges, technically, and from a development point of view in getting this thing ready to produce?

Grogan: Well, there were many. <laughter> There were many, but the one that is really quite vivid in my mind had to do with the plan we had for thin film heads. We began work on thin film heads in the ’60s at Yorktown. They were single-turn heads, and we were accustomed to 20 and 30 turn heads. I worked in Yorktown for a short time to design a write driver for a single-turn head, which was a one-amp write driver. <laughter> The transistors were in oil for cooling. <laughter> And that demonstrated the notion that a thin film head could write something, and read something back. But it was clear in a commercial application that we needed higher turn heads, and with 3380, we contemplated thin film disks with very high corecivity. And so we thought this was going to be simple, because we’d have a dual-element head. A write element that had a very wide gap, low-throat height, large pole tips that could just write on cardboard, if necessary, and a read element whose structure was the antithesis of that thin pole tips, thin gap, so forth, so that we would have the best of both worlds -- an efficient writer, and a very efficient reader. And it was quite stunning to us to find that the write element coupled into the read element, as you passed over data, and created an artifact in the read data, whose frequency spectra was identical to the data being read, and therefore, there weren’t any tricks that we knew how to filter out this artifact without filtering out the primary read circuit.

Lucke: Uh huh.

Whitney: That’s right, the data.

Grogan: So that was a very disconcerting moment. Attempts were made to have dual-layer write elements, 16 turns, two layers of eight element, eight turn heads. Those were very difficult to fabricate, and I don’t recall us getting any to work satisfactorily. And we quickly arrived at the point where it was clear if we were going to use a thin film head it had to be a single-element head, whose internal geometry was sufficient for writing and reading. And so we built such a head. An eight-turn head that traded off pole tips and so forth. And much to our horror we discovered that at 15,000 bits per inch the leading and trailing inches of the head coupled into the read signal, and spectrally we had a similar problem. And it was clear that we could not -- this was in about 1978 -- that a single-element thin film head would not work at 15,000 bits per inch. So there was much adieu about moving the inside track out to get the bit density lower, to avoid this problem, and raising the track density to recover the capacity. And all of those solutions were really impractical. In the 1970s, Peter Franchech at Yorktown lab had invented two new codes that came to be known as 27RLL, and 17RLL, and these codes offered flux change efficiency in exchange for clock speed. And in the case of 27RLL, it took 66 percent of the bit density would result in flux changes, but the data bit to clock expansion ratio was two-to-one, so therefore, at three megabytes per second, we were facing a clock running at 48 megabits per second, which we all thought was -- I wouldn’t go so far as to say impossible, <laughter> -- but as close to impossible as anything that we had ever considered -- as a matter of fact our oscilloscopes really weren't -- the Techtronics offerings in those days, you couldn't really display a pulse operating at that rate that was only 10 nanoseconds wide. To make a very long story short, a guy in my group designed an encoder, a 27 encoder that was sort of a prehistoric version of the true encoder, and allowed some encoding to demonstrate the principal of Franchech's invention. And Paul Hodges and John Eggenberger, came up with a very minor modification to the algorithm that made it practical to put a 27 encoder in the 3380. We had so little confidence in the encoder because it had a quirk that if it ever got one bit out of step, it propagated an error all the way to the end of the record. So everything that was written was passed through the encoder and onto the write driver, and the output of the encoder went back through a decoder and compared it to the input to the encoder to be sure we had not made an error. And so by dropping the flux change density from 15,000 to 10,000 we were able to get pole tips that were large enough to (a) not interfere; and, (b) that were able to write. And I remember the first head of that design at that density that actually worked, and that was one of those great moments in science. <laughter> Because quite frankly, we thought we were really screwed on this one.

Porter: And your technology was followed and copied by innumerable people afterwards, wasn't it?

Grogan: Yes, it was. But without the flux change efficient codes we would not have got there. Franchech later won an enormous corporate award for the development of the invention.

Porter: Okay, and other things that made this all possible. The disks. There was a transition from those old oxide things.

Grogan: Well, there was a transition to and back. <laughter> The original -- this is where my memory is fuzzy, but as I recall, the 18-inch disk with the four actuators, horizontally opposed and so forth...

Coolures: That went early.

Grogan: That was very early, and it was a plated disk. And I think it was nickel cobalt, and it had difficulties in terms of corrosion. And I think there were various overcoats of chrome and what have you. And then the 14-inch disk, again, this was a size-power consideration. The 14-inch disks were also plated. Originally, the substrates were bronze, as I recall, to get the smoothest possible substrate. But we had great difficulty flying on that. And Andy Gooday and I worked on sort of a skunk works project to see if we couldn't make a brown disk, an oxide disk, thin enough to support these densities, and we could. So we had both plated disks and brown disks in HDAs, and I don't recall now what the decision -- you may recall, Chris, what the decision was based on.

Coolures: Well, I think...

Grogan: Going back to the brown disk.

Coolures: I think one worked, and one didn't. <laughter> But it was -- well, I mean, giving John and Gooday a lot of credit, because they were -- in the skunk works they were really pushing the particulate disk, and actually demonstrated that it would perform, not only mechanically, but also it would perform from a start/stop point of view.

Grogan: We actually built some HDAs that everybody thought were plated disks, and we blew them in the drive.

Lucke: Well, I think the mechanical side was key, because we were having a lot of trouble with the start/stop life on the thin film disk. And the particulate disk looked much better.

Warner: Well, there's both. The start/stop life as well as random crashes that occurred.

Grogan: Yes.

Warner: We would be in product test lab, and then we'd get a call, "Okay, we crashed." And we'd go there and the heads are ripped off, and the disks are torn up, and so it would -- we'd just get to the 11th hour, and we'd think, "Oh, we're almost there." We're through product test, and then the thing would come apart on us.

Grogan: Yes, I remember that.

Warner: And we'd try different kinds of lubes, and different lube thicknesses, and different glide height things. You can't see this, but the disks on this drive are twice the thickness of the previous generations of oxide disks. You have to remember this, IBM is planning on 3380, has a big factory all set up to do film disks of this non-standard thickness. And the reason for the non-standard thickness was because of disk flutter within this -- within the HDA. And then to make this change at the last minute to an oxide, you know, it was like monumental. It's like, you had to pull the baby out and start all over again. It was a major, major change for the company.

Porter: So the program did start out then with the oxide disk?

Warner: It started off with film disks.

Lucke: Yeah.

Warner: From the beginning, film head and film disks...

Porter: I mean, when it actually shipped.

Warner: When it shipped, it actually shipped with oxide.

Grogan: With an oxide disk, yeah.

Warner: The delay that was involved...

Coolures: And it was announced with an oxide disk.

Warner: Yeah.

Coolures: Through the actual formal testing after many failures with the film disk, then we went with the oxide disk. So all the final testing was done with that.

Whitney: Well, all the pre-announcement testing was done on an oxide disk. Not all of it, but the testing that qualified the product for announcement.

Coolures: Yeah.

Whitney: Was on oxide disk.

Porter: Well, how about the rest of this drive? There were actuators, there were motors, there was ways of moding errors and outs, and all the right things. Where were the real problems?

Warner: Everywhere.

Lucke: I could tell you a few of the problems. <laughter>

Coolures: Well, why don't you talk about the actuator, because that's a unique design in the sense that the ways typically were, you know, fixed on the HDA, and the carriage had the bearings, and this was a complete reversal on that.

Lucke: If you look at all the previous programs we had, you basically had bearings that were riding on a cylinder. And so the actuator had bearings on the bottom, and it would go on a round cylinder. The concept here, because we were really worried about getting as much symmetry as possible, was to make the actuator as symmetrical as possible. And so the decision was made to actually put a plasma spray, and actually spray it on what we called the bird and this was the cage that had the bearings in it. And we actually had a guy named Kay Ikeowa, who actually had worked in plasma spray in his previous life before he came to IBM. And I would say this was one of the most difficult manufacturing problems we ever had, because if you get any debris that comes off of this, it is inside the HDD, and that debris gets to the head disk interface, and you get a crash. And we spent-- the designer was Pat Merrill -- and he spent probably a year working on this design. Because you have to basically spray it, and then you have to grind it. And then you have to get it as flat as possible so that you don't have any problems, and you got to get it as smooth as possible so you don't generate debris. And we had to worry about the hardness of the bearings, the hardness of this, and make sure they weren't mismatched, because if one was harder than the other, you had problem with the debris being generated. So it was quite a process. In fact, the manufacturing engineer, Dick Trammell, some of you may remember, came into me one day and he said, "I'll make anything you want, but this design is unmanufacturable." <laughter> And so a year later, I agreed with him, but he'd already figured out how to do it. <laughter> But it was quite a unique design, and it did help us in some of the track misregistration problems, because we were worried about the dynamics of the actuator, the dynamics of the disk flutter, you know, the tilting, the thermal effects. All the things that were associated with this actuator were quite exciting.

Coolures: Yeah, I'd say that was the toughest part, and then the coil itself was a problem, because it's a structural member. And that had its learning curve.

Lucke: Again, just like with the actuator, this particular design, normally we wound it on an integral structure. This was a first design, where we basically made it free-standing. And getting this design, and getting it coated, and make sure it didn't fall apart was very interesting. So there was a lot of unique new ideas that went into the 3380 that hadn't been on any previous programs.

Grogan: Well, I mean, the result was the moving mass was very low, right?

Lucke: Yeah.

<overlapping conversation>

Grogan: Because you weren't moving the bearings.

Warner: You weren't moving-- it was symmetrical.

Grogan: It was symmetrical, and it had very low mass.

Lucke: Right.

Whitney: And by having the coil be the structure, you also were able to reduce the gap in the voice coil.

Lucke: In the voice coil motor, right.

Warner: And so you were more efficient from a magnetic standpoint.

Lucke: So that was one of the problems that we had. A second problem that we had was, if you look in this design here, we had this very large bearing spinning at 3600rpm, which is quite fast for this size bearing. And we're worried about very, very little run-out and all the problems. So when we first looked at this design, we'd had problems with noise in previous designs. And so we thought, "Well, why don't we put a quiet, plastic retainer in this bearing?" That was going to be the solution to all of the cage instability problems that we had. And it turned out that this actually took up more space that you could've had grease in there, and the important parameter was the grease. Not the plastic versus the steel ribbons. You can see the design that we ended up with, we went back to the normal steel ribbon retainer, because you could actually pack more grease around the outside of the retainer. And that solved most of the noise problems. But...

Warner: Now to me this whole thing is kind of interesting, because what drives a lot of this is that this was-- the heads were start/stop in contact. And so you had, you know, what is there, 32 heads inside this thing, that you have to start and spin up very quickly. So that means that you have to have a big motor. Do you remember, Chris, what the horsepower was?

Coolures: Oh, I don't remember the horsepower of the motor, no.

Warner: It was like a 3/4 or one horse motor, or something. It was a big motor. And this was the hub that the belt went around. So that big motor had to have big tension on it, in order to transfer enough energy in to spin these thick disks with all these heads sticking on it. So I mean, one thing led to the next, to the next, to the next. That kind of, you know, complicated. If we hadn't been start/stop in contact and weren't critical to the spin-up speed, then you could've done this with a much smaller motor, and a smaller bearing and so forth.

Grogan: But you know, Jim, you should comment. I mean, I was a Double E, so I was sort of a spectator as far as this bearing was concerned. But my impression was that you guys were really into the world of bearing science really early. I mean, you went places that nobody else had ever gone before in terms of...

Coolures: That size bearing, yeah.

Lucke: I actually spent time at the vendor in Germany working with their scientific department, because they had never had people come back and "say this bearing isn't good enough". And I remember several trips. In fact, I even remember one trip where I actually brought back sample bearings and I got stopped at customs, because they wanted to know what I was carrying that was full of steel. But it was quite a lot of work working with them to get the right amount of grease, type of grease. We went through several different types of grease. Went through all kinds of work with theoretical equations on age and stability. It turned out the secret was to get rid of this, and just make enough space to get enough grease in there. <laughter> In fact, I think, Jack, you spent time later with Chuck Lanzi, where he went actually went in on some of the early drives and actually re-greased some of these drives in the field.

Grogan: Yep.

Coolures: But you know, like any product, I mean, even after going through all that, then we started having some problems later. And we found out that the vendor had actually just modified his process a little. He didn't know what he had done either. And that took a long time to sort all that out. I can remember a lot of people spent a lot of time in...

Porter: Exactly. If you went through all the development activities to get this thing ready for final introduction. At some point, the management involved was confident enough that it was ready to go. That you actually then on June 11th of 1980, actually made the announcement of the product, and announced the delivery date. I don't remember what date you announced.

Coolures: It was first quarter of '81.

Porter: The first quarter, yeah. But it didn't actually ship until October of '81.

Grogan: Right.

Coolures: Right.

Porter: Which was an embarrassing delay.

Coolures: Right.

Porter: So we'll need to, I think, of necessity, discuss what kept all of that beautiful development work from really being ready for production at that point in time.

Coolures: Well, you want to answer that? <laughter> Well, no, what happened was we continually do a lot more testing even after we've announced the product, before we do first customer ship. And one of the things when we went to the particulate disk, and we were concerned about start/stop and stiction characteristics of the head on the disk. There was another lubricant that we had tested before, and it had excellent properties on start/stop and stiction. Much better than the lubricant we'd used on the 3350, and that. So we were in the process of testing that, and then we started to get some failures after like 180 days. So that's when we decided that we had to really find out what the problem was, and we had to do the testing to qualify that. So we went back to the old lubricant and still had to retest the product to make sure that we didn't create another problem, and that was the reason that we delayed the shipment.

Porter: Which took a lot more time and you had promised to the people you reported to, then, huh?

Coolures: Well, I mean, when we announced the product, we didn't have the problem. And the problem didn't show up until after announcement. And so then to be able to recover, you can't just make a change like that without going through and retesting everything and make sure that you have everything fixed.

Grogan: The difficulty was the gestation time of the problem was so very, very long. It was in this four to six month kind of thing. And so in order to know that you did not have the problem with a possible fix, you had to then test a long number of drives for that length of time. And so that is what drove the enormity of the delay, right?

Lucke: We spent quite a lot of time trying to figure out accelerated testing. And it turned out accelerated testing never really represented reality. You had to actually run a big sample for a long period of time.

Coolures: Right, yeah. Because we always did a...

Porter: Hold it there. Given all the difficulties in making it happen, the initial versions of 3380 were actually shipped in October of ’81. And during 1981, the estimates I made when I was doing the market studies, was they shipped 400 HDAs, which I guess would’ve been the equivalent then to about maybe 200 boxes, something like that.

Coolures: Yeah.

Porter: And it increased rapidly. In 1982 we estimated 14,000 HDAs for 3380 shipped. Then in 1983 51,000. And the next year in 1984, 74,000. in other words it had become a rather successful product. And looking at those prices of $100,000 per box, that was a pretty good revenue source for IBM. And it created, of course, a lot of difficulty for those nasty competitors out there selling those older 3350 drives. So it really changed the scene for IBM in a business sense. And where did IBM make all these drives?

Coolures: Well, I mean, we...

Grogan: Everyplace! <laughter>

Coolures: In San Jose, and Mainz, Germany. In Vemecarti, Italy, in Brazil, and in Fujisawa, Japan. And that was -- and actually it was amazing how having these products built in all these different locations, and what I used to use is the whatever one was doing the best job, that was the target for the others. And that really gave us some advantage in proving the process and getting better quality.

Lucke: It was also interesting that actually putting that technology into other places was quite an exercise from a manufacturing point of view. San Jose and Mainz were the two main locations. But we had to put it in Brazil, for example, and Brazil didn’t really have a lot of local vendors that really knew some of the technology, and to produce in Brazil, you had to put some of that technology into Brazil. And I can remember one experience one of the times I was down there putting in a clean room, that they took me out to the vendor that actually made the clamp for the disk. And it turned out that vendor made rifles. And I went into his facility, and it was like my garage. It was just everything all over the floor. Rifles, gun barrels they were boring. It was just filthy! And then he says, “Let me show you where we’re making the clamp.” And we walked around the corner, and he had constructed a complete clean room that made just the clamps. And it was like the shining light in this room, and then the whole rest of the place was making gun barrels.

Coolures: <laughs> Right.

Lucke: And that's what they did, is they actually tried to bring up vendors, transfer the technology to them, and they actually monitored the cleanliness level in the room, and everything else. And that was true in a lot of places. And it was quite interesting. And we did the same thing in Vemicarti. We built a clean room there also, as well as Sumare.

Porter: Well, while all that was going on, the development people were being told you have to do the next version, right?

Lucke: Yes.

Porter: And in 1985, in February of '85, IBM announced the 3380E models. Which had twice the capacity per HDA. That must've been an interesting challenge.

Warner: I have a little story about that. Sometimes during the product planning we would get together at some offsite place. And we went down to the La Playa Hotel, and we had probably 50 or so people in the room. And they were people from the technology areas, and heads and media channel, servo system, the box manufacturing people, all in the room. And this was being done right in the throes of where we're trying to ramp up the 3380, and we're having all these problems that Jim's talking about, and Jack talked about. And everybody you know, the plan was to double the capacity. And the team said, "No. No way. We can't do it." And so we talked and we said, "Well, what can we do?" And so the whole team signed up for one-and-a-half times. And we said what we would do is go back, do our work, look for doing two times, and then we would meet again in whatever it was, six months or something like that. So we walked out, very disappointed, with only one-and-a-half times capacity. And at the next meeting, everybody had done their work and came back and said, "Oh, we'll do two times." But I believe if you said you had to do two times, that team would've proved you couldn't do it. But if you said, "We'll do one-and-a-half, and then we'll give you a chance to prove it," it came back and turned into a double density.

Porter: Which you had several years to do it as it turned out.

Warner: Uh.. yes, but there were a lot of major problems. Mostly doing-- because it's all done with track density, mostly servo and read channel, and heads media, I think were the technologies.

Grogan: Well, there were -- because of the problems geometrically with the single-element, thin film head, there wasn't much room to increase the BPI. So we had to do it 80% with a TPI. And what I recall was in order to reduce the TMR at that higher TPI, I believe we lowered the average access time. So that there would be more time to settle. Is that-- yeah.

<overlapping conversation>

Warner: Jack, you should say you're using some acronyms that people might not know.

<overlapping conversation>

Porter: Well, what we're talking about...

Grogan: What acronyms?

Porter: Tracks per inch went from 800 on the original model up to 1386 on this model. And the bits per inch, the BPI, went from 15240 up to 16200. This, every seek time actually stayed only one millisecond.

Coolures: One millisecond, but that was...

Warner: That was the settling time.

Grogan: That was huge! Because the settling, that's where the track mis-registration, I mean, where we really paid quite a price. And one additional millisecond reduced that dramatically. You'll notice that when we went to 3380K, that we actually dropped the bit density back, because we were living in pretty thin air with a single element head at 16000 BPI. So we had to drop -- we actually dropped the bit density back, and took it all in TPI for the 3380K.

Porter: Well, I think the specification change that IBM management probably liked the best was that the sale price went from $101,000 for the A model, which had two HDAs, when you came out with the E system, it went up to $134,000. So that's the kind of specification change most managements would like to see, isn't it? But you double the capacity for the user, so you gave them a deal.

Coolures: Sure.

Whitney: I think it's useful to point out at the time we had to delay the first customer shipment, not only were there myriad technical problems, which have been described, but it created a serious business problem for the corporation. And the reason being that we had sold these products, and customers were waiting for them. Customers didn't know for sure then what their plans were going to be. So there was fear that competitive choices would be made instead of the 3380. And worst of all, as we had planned to scale up of production of the 3380, we had planned the scale down of the 3350. And so we didn't have any products to send to our customers for a period of time.

Coolures: Right.

Porter: I can tell you that the folks in Colorado at Storage Technology became a killer on that. With their double-capacity version of the 3350, they made a real killing on that delay.

Coolures: That's right.

Whitney: But I think the comments made by everyone reflect what to me was one of the most important aspects of IBM during my career, and that was the integrity of the corporation, unwillingness to ship a product about which we did not have complete confidence, and we took the risks and took the consequential business impact to make sure that the product we finally shipped was going to be right.

Grogan: Jim, you're probably aware of this anecdote, but a colleague of mine, Jack Schwartz, actually built a double-density 3350 within IBM, and a conscious decision was made to not pursue that product, in favor of the 3380. And I don't remember the circumstances, but one of the difficulties that was sort of frowned upon at the time was, as you know, STK, when they installed the 3350, I mean, once they set it up on the floor, they formatted the drive in the customer's office to account for any shipping tilting, or any weirdness that went on. These were count key data format machines. And the count fields were all written in the factory. And it was one dedicated head for a servo. So if anything tilted, at those high track densities you were in trouble. So STK saw no trouble with formatting in the field. The customers, as it turned out, so no problem with formatting in the field. But I could be recalling incorrectly, but I think was a factor in our decision. To not go forward to 3350. I mean, we had them running, but they were running...

Coolures: No, that's right.

Porter: A tactical decision.

Coolures: Uhm hm.

Grogan: Yeah.

Whitney: Well, I think that we felt that we were not moving the technology ahead, and therefore were not enhancing our competitiveness with a simple double-- or with a double density 3350.

Coolures: And if we did that, it would've diluted the effort on the 3380.

Grogan: Yes, absolutely.

Coolures: And delayed it even more, and so it's where do you put the resource?

Porter: Okay, but the 3380 progressed. It went to the double capacity version of the 3380E in '85.

Grogan: Uhm hm.

Porter: Only two years later in September of '87, you announced the 3380K series, which of course, took the capacity three times the original up to 3.7 gigabytes. And this became the basic product, because it within a year or two replaced the E model. And became the basic product you were shipping at that time. Same kind of difficulties in going up to three times the original density?

Lucke: I think from the mechanical aspect, you know, we talked quite a bit about the technology, but every time we progress with the technology, we had to look at, for example, spacing of the heads to the disk. 'Cause you could read better if you float closer. And if you fly closer, the asperities in the disk cause you problem. So we spent a lot of time in the mechanical interfaces making sure that we could actually support the technology. Not only from the mechanics of flying the head, but also we had to worry about this track mis-registration, and minimizing all of the different parameters, non-repeatable run-out, thermal tilt. All these things we had to keep squeezing them down so we could support the next track density that we needed.

Porter: Well, these things did cause some problems. Let me quote you an article from the *Computer Systems News* of September 25, 1989, written by an old friend of mine, Hesh Wiener. And here's a paragraph. "A new generation of mainframe disk drives would not throw the industry into turmoil if everything were going well for IBM. But in the past two years the company has suffered unusual and embarrassing problems in its mainframe disk drive business. The most dramatic errors was the deployment of thousands of 3380K disks built with spindle bearings prone to failure. The bearing problem affects every 3380K shipped by IBM from the first customer deliveries in the first quarter of ’87, through deliveries made in the summer of ’88.” Interesting problem.

Lucke: Yes.

Porter: You got some pretty good press releases.

<laughter>

Coolures: And that’s the bearing.

Lucke: I remember going to a customer, I won’t name, that had one of these problems. And I walked into the person that was running the data center. And I went through about a half-hour explanation of why we had problems, what we were doing about it, what we were doing to replace things. And they looked at this bearing, and the whole time I’m talking to him, they’re staring at this bearing, ’cause it was one like this. And at the end of the discussion they said, “Okay, can I keep the bearing?” <laughter>

Coolures: ’Cause he probably wanted to show his boss what the problem was.

Lucke: They wanted to take this back and show it, and say, “This is what’s been causing all the problem.” I think it ended up a paperweight in that person’s office.

Grogan: Yeah, it’s a bearing chirp.

Lucke: Bearing chirp.

Coolures: You could hear the chirp of the bearing.

Lucke: It was an audible noise. You could see a little bit on the mis-registration, but it was the audible noise that got people concerned.

Grogan: I think, again, being a Double E, <laughter>, being on the outside of this...

Coolures: We forgive ya!

Grogan: But what was to me I found this rather incredible. But mechanical team, Jim and Chris and others found a way to replace this bearing in the field. We actually removed the bearing in the HDA in situ, and replace the bearing. I don't remember why I got -- and many, many others were-- conscripted to go the field and replace bearings...

Lucke: 'Cause you're a good guy. <laughter>

Grogan: We were weak-minded or something. But that was -- that innovation -- I mean, do you remember how that came about? I have no recollection of how it came about.

Warner: Because you could get the pre-load.

Lucke: Yeah, what you actually had to do is you actually had to pull it out from this side, which means you had to...

Grogan: No, but who devised the idea?

Lucke: I don't remember exactly who the person was that devised it.

Warner: It was a bearing puller.

Lucke: It was a bearing puller.

Grogan: But it seems like an impossible, you know...

Lucke: It's like changing your engine while you're driving 60 miles an hour. <laughter> Except for we allowed you to slow down on this one.

<overlapping conversation>

Lucke: It was kind of the court of last resort, though.

Porter: Which is pretty messy work for a computer room.

Lucke: Yes.

Porter: But the problem was painfully and painstakingly solved by IBM, and the program continued.

Lucke: In fact, it turned out the only thing that helped you in our subsequent products was the fact that we had to get this thing smaller. So as you see with time, we kept getting the bearings smaller and smaller and we found we could spin it, because everybody wanted to spin even faster than the 3600rpm. So in the future because of latencies they wanted to get as much speed as possible, and we said, “Look, this guy is killing us, don’t run it any faster,” but eventually the solution is to get much smaller bearings. You know, it has so much inertia to be spinning around.

Porter: Were there any other messy problems along the way?

Lucke: Well, I think there’s one that I mentioned that goes with this that maybe Mike wants to talk about, too, is you think about our solution to this problem is putting in grease. But it turns out that this bearing is actually inside the HDD, which means you have to worry about whether or not that grease contamination can get in the head and the disk. And you say, “Gee, a little extra grease helps to lube,” but that’s not exactly what happens. And so it actually turns out that part of this design we had to design an elaborate labyrinth so that you actually made the flow from here back into the HDD very torturous so that nothing can get back in. So if you looked at the cutaway version of this, you’ll see this funny labyrinth.

Warner: Yeah, like this.

Lucke: Yeah, it looks like a gear tooth, and you had to basically go through that path to get inside of it.

Warner: And the air’s trying to pump out. And Chris you might-- the airflow system had a few issues along the way, too.

Coolures: Well, it was an open-loop system, so you had a blower, and you had an absolute filter. And actually we -- with all these machines cycling as much air as we did, we really had -- in the computer room, a very clean room. See? And the life of the -- and we had to design a filter that would last the life of the machine. We didn’t want to replace filters and that. We did have some problems in the field. One was a contaminant was even getting through the filters. It would accumulate on the heads and cause crashes, and we had a -- that was a real headache for us for a long time until we found out what the problem was. And what it was was that our customers were using tributyl-10 oxide that was supposed to only be used in cooling towers in their clean room in the air-conditioning system. Because people would complain about there’s a musty smell in the air. And this one vendor they were using it in the clean room, and that tributyl-10 oxide would actually accumulate on the head, and eventually cause a crash. And that was a big headache in the after...

Warner: We'd go right -- by you to mounts. We'd go through the filter, and as the head's flying over the disk, it'd taper flat, so it -- as the air comes through, it's like a perfect little catcher. And it would accumulate little bits of tin on the surface until finally that would react or gum up the lubricant, and it wouldn't lubricate anymore, and out it would go.

Porter: Well, despite all these problems, the 3380 series had an exceptional product life. Actually from first production in '81, the last production wasn't until 1991. We're talking about 11 years of actual production and sales. And there aren't too many disk drive programs that have ever had that kind of history. So uh.. overall, despite all the nasty problems we've been talking about, it was an exceptional product life, and made a lot of money at the bottom line for IBM, obviously.

Whitney: Right, well, I think one thing to keep in mind is through out all those models and the progression of the technology to the user, the interface looked the same. He was getting a data rate that was the same as that which he started out with. And he was getting record lengths that were exactly what he started out with. So it minimized the conversion or transitional problems that new products every three years sometimes introduce.

Porter: Well, in '89, a couple of years before the end of 3380 actual production, IBM did introduce the successor product, the 3390, which of course, was advertised again as using a several times smaller footprint in the computer room. A much bigger price per megabyte. And all of the good things that went along with it. So IBM continued with that advance. Went to a smaller disk, 10.8 inch, is that what you called it? Were any of you also progressing on to the 3390 program after 3380?

Grogan: I did.

Whitney: And I did, too.

Coolures: Yes.

Warner: I was at the launch.

Whitney: Yeah, early on.

Warner: I was at the group that started off with the design.

Porter: Was it as satisfying an experience to get it going? And with as much long nights and grief?

Warner: Well, I left IBM so I wasn't there before it shipped. And I quite frankly wanted IBM to make smaller form factor disk drives.

Porter: It was first shipped in 1989, yeah.

Warner: Yeah, and I left IBM in '85. I campaigned very hard to make smaller drives in arrays. You know, RAID arrays. But this was, I guess, the last of the big iron, the 3390.

Lucke: We had the 5-1/4.

Grogan: I think Mike is being quite kind, describing. Because my recollection is that there was a rather extraordinary controversy surrounding the disk size to be used on this drive. I mean, it was almost religious. I mean, there were various churches and congregations supporting different views on what the disk size could be. And I think Ric Davis was the project manager at that time.

Warner: Right.

Grogan: And he was trying to sort out technical fact from emotion for several years over what was the disk size to be. And I think that -- I mean, there were a million factors, power being one, smaller bearings. But there was also compatibility insofar as record lengths and so forth. And so you ended up when you took all of the requirements that were actual peripheral to the design, you ended up with not many choices about what you could do. If you wanted to raise the RPM, and you had same limitation on data rate, then suddenly you got locked into bit density, which translated into the same number of record-length compatibility and so on and so forth. But it was quite complicated.

Lucke: The bit density, my view, was one of the bigger guys. Because what happens is, if you want to keep the bit density, and you get a smaller form factor, you know, you have a problem. So they wanted to keep the size big enough so we didn't have to spend the thing faster. And 'cause you got the radius effect on the bigger disk.

Porter: When they went to the 3390, they went to 27,940 bits per inch, which was up from 15,000 on the last generation of 3380.

Lucke: Right.

Porter: So yes, there was a very major improvement in bits per inch.

Whitney: That was a thin film.

Coolures: But again, it's all driven by compatibility, so you move the data from 3380 to the 3390.

Grogan: We had learned a great deal about -- in thin film heads about -- I mean, what really mattered in the structure, the materials, the process and so forth. And so we just made more and more efficient single-element thin film heads. And as I recall, Dick -- and I think you were one of the people helping us to see the light -- that improved latency was a necessity. Because latency went right to the bottom line in any model that we ever...

Whitney: We were getting lots of input from the people in Poughkeepsie, who were manufacturing these big mainframes, and trying to accommodate the transaction driven systems, and for transaction driven systems, latency was crucial. We wanted to bring it to zero.

Porter: Well, you went from the revolutions per minute in the 3380 from 3600 to 4260. A modest increase, but you know, a lot of reach.

Lucke: But that modest increase drove a lot of problems relative to bearings, again. Whenever we thought about increasing the speed, it brought up new bearing questions.

Grogan: Well, the total power to spin a disk varied something like -- by the cube of the RPM change. So that modest increase, if we had tried to do it with a 14-inch disk, the power would've been incredible. And so I think the trade-off was a smaller disk and higher RPM, and I don't recall. But probably...

Coolures: No, I think that was it. We looked at what's the smallest disk, and we could still make it so it's compatible to migrate the files. And that kind of where the design point had to be set.

Whitney: Yeah, there was a kind of an area in there where you could make that work.

Porter: Of course, things kept getting smaller. After the 3390, IBM started using for a few years 5-1/4 inch drives before they moved to 3-1/2. Most of the competition went directly to the 3-1/2, as I recall.

IBM did a stage there where they used storage subsystems using 5-1/4 inch hard drives for servers. So it kept going down. And just to project into the future, of course it did go down to that diameter, and who knows where it'll go afterward. But it was all a matter of what could be done the most cost effectively with power, heat, and all those other things that the computer room is sensitive to.

Grogan: Well, within IBM, the 5-1/4 development in Hursley used rotor reactuators. And I think San Jose's product, which was called Sawmill...

Lucke: Sawmill was linear.

Grogan: Was a linear motor, and I think it was the last linear motor ever put on a small...

Coolures: Yes, it was, yeah.

Grogan: ...disk drive.

Porter: Probably, yeah.

Grogan: Yeah.

Porter: Because the rotary actuators which moved this way rather than like this, of course, really replaced everything. There is no disk drive made today with a linear actuator. They all use rotary.

Coolures: Oh, no, especially as you get smaller.

Grogan: Yeah, everything with rotary actuators had sector servos. And the linear motor device had a dedicated servo, which was kind of the San Jose ethos, I guess you would call it of those days.

Lucke: I think one of the key things that sort of helped the transition to the smaller form factor was this, what we call, meantime to failure. We spent a lot of time making sure that these HDDs lasted a long time. Seven years, plus, was the minimum requirement. And there was a lot of concern of the reliability of some of the small form factors. And as time went on, they became more reliable, then we started using more RAID systems, and pretty soon you were able to put what we had considered low-cost drives into these very expensive computers. Originally, we didn't think we could do that. And we had to put more reliability into the original of the design. So as soon as you got over that hurdle, it opened up the door for all of the product we have now, which all have gigantic arrays of small form factor, low-cost drives.

Coolures: And I think also we were always into the linear actuator in San Jose, basically, and Hursley had the rotary actuator. But their approach to a rotary actuator was not compatible to making a smaller, smaller file. Because it was still a big motor, and when they went to the -- I don't know what you call them now -- but the rotary, but where the magnet and the coil is in the flat form factor, then that really opened up a whole new line of products you could develop that way.

Warner: Well, the cost point changed dramatically. I mean, this bearing, I'm sure, cost more than a whole casting/bearing/motor assembly for a 3-1/2 inch drive.

Lucke: It does -- did.

Grogan: It's by the pound. If you have the volume, then it's by the pound.

Porter: And the quantities of the drive have gone way up, and the price per -- we don't even talk about price per megabyte anymore, nowadays, we just talk about price per gigabyte. And you can buy drives today in certain resale locations at 50-cents per gigabyte, which is a long way. And you guys made a significant contribution to that as a step along the road. So perhaps to wind up today, I'd wonder if we could get each of you to give us a little bit of your thoughts as you look back on how you think about your participation in this whole thing, and what you feel was the result of all of this. Mike, why don't you give us a few thoughts?

Warner: Well, for the last ten years, I've been working in semi-conductor packaging, which I think is making a difference. But I look back and I think of my career at IBM, and the 3380 was like ten years of that time, that we really set the technology, set the trends that would make a big difference in the world. I mean, storage is an enormous part of a computer system. We helped to develop the technologies that would be used for miniaturization later on. And the heads of media technology, the start/stop in contact, the servo, the read channels, all of those things were critical to a vital part of the industry. So you know, I can look back, I can tell my kids -- I don't have any grandkids yet, hopefully, I'll have some -- but that about a time that I think that I worked in a company, in an industry that made a difference to the world. So I'm very proud of it. I look back as a great time. These are the kind of guys I worked with. They're smart, they're down-to-earth, they make things happen. They'll work through the night if that's what needed to get the job out. So I look at it as a fantastic experience!

Porter: Good! Jim?

Lucke: I think that the biggest contribution that I made was really in what we called this track mis-registration. I spent from the day I think I joined IBM trying to figure out ways to basically minimize that offset between the head and the track. And I even spent a lot of time developing models to statistically look at whether or not, you know, the probability of being off-track a certain amount was greater than what Jack and his channel could tolerate. And so we spent a lot of time. In fact, I remember in the 3350, I was told, "We can't do it. It's too much offset." And I actually proved to them that statistically you were okay, and we actually did the 3350 with that analysis. And so I spent a lot of time looking at all the things that contributed to this offset. Whether it was vibration, whether it was mechanical, tilting, non-repeatable run-out, all those kind of things affected it. And I think that parameter, along with the technology was the key driver to get us to go to the higher density products as we went on. So.

Porter: Good! Jack?

Grogan: Well, I was never much of a visionary throughout my IBM career. When I began, the drive of choice was the 305 RAMAC, the original drive, and five megabytes. And at the time I couldn't quite imagine what anybody would even need five megabytes for. It seemed like more data than is possible to use. In 1960, I was sent with another engineer to the Coast Guard Academy at New London, because they thought they needed 10 megabytes, and we all knew nobody needed 10 megabytes. And we came back, and they actually needed 20. And my manager told Bill Benz and I, "I knew I shouldn't have sent you guys." So at that point I thought, "You know, there may be something to this business." And I had the opportunity to work as IBM went to 3-1/2 inch drives, when capacities got to 10 gigabytes, 20, 40, and so forth. And so I have been completely amazed. And my grandchildren now, two of them, because they store digital pictures, actually we got them 60 gigabyte drives for their laptops for Christmas. And so I just can't imagine where this whole notion of storage is going. How it'll affect the human condition. I mean, I know what it does in medicine and other places. It's been a rather extraordinary journey.

Porter: Chris?

Coolures: Well, I kind of feel the same way. I think, you know, working for IBM during that period, I started in '60 and retired at the end of '91, and worked on many products, including a high-speed printer. And one thing I remember is that there was always that horizon in the disk development that you could look out ten years, and you knew, or you had confidence you could make a product to still improve performance and cost-per-megabyte and all that. And it was always that ten-year horizon. And even when I retired, you could see that the files were getting smaller. But it's amazing how far it's gone now. I mean, it's just beyond belief. But I think the net of it is that the thing I remember most is the people I worked with. And that's the thing that I miss from retirement, is the people I worked with. And it's been a great career!

Porter: Good! Dick.

Whitney: Well, I share the same views. It's a privilege. As a civil engineer, it's a privilege to work with these guys. <laughter> But I think the thing that has changed from my point of view is that starting out, when I first came to San Jose, and I think throughout most of that time, we tended to look on disk storage as the product. I mean, that was our entire focus. Whereas, now, when all of us have our desktop, or our laptop, a computer more powerful than the 3033 that was the top mainframe at the time the 3380 came into being. We now can step back and look at these elements -- the disk drives, the memory, the displays, the input -- you can look at that as a system. And I think that now disk development is really done from a systems perspective. Much more than perhaps it was at the time. Even though the guys from Poughkeepsie were trying to beat us up and tell us that we need -- and we didn't beat back hard enough, tell them, "We need a serial channel now!" <laughter>

Lucke: You know, one thing you add to that is...

Porter: We really are out of time, I think. So I just want to thank each of you for being a part of a fabulous discussion over how some significant things were made to happen, that would not have happened without all that hard work. Thank you very much, gentlemen.

Warner: Thank you.

Whitney: Okay, go ahead.

Lucke: No, I was just gonna say, you know, we spend most of our time today figuring out the compatibility of our product with all the different operating systems. So that testing is probably 90 percent of all the testing we do, because everything now is a sub-assembly. The HDD that we worried about is a sub-assembly we buy, and it mixes into this overall subsystem. So it's a different world!

Whitney: Yeah, I can remember -- we're all flying now I guess...

Porter: Well, he's just about finished, I think. Oh, you've got three minutes, go ahead.

Whitney: Well, the idea of is the disk drive an industry, or is it a business independent? If you remember that period when IBM when we were -- who should we spin off, and what should we do? In retrospect, it was not. I mean, it was. But it wasn't the business that IBM was in. And IBM was really in the systems business.

Warner: That's right, yeah.

Whitney: Then the other thing I remember, which some of you also remember, was the difficulty of convincing some of the executives of the company why we were gonna -- or how we were gonna use all this capacity. Because the mindset was the data arrived, or you know, was input to the storage system as a consequence of keystrokes. And people were calculating how many keypunch operators is it going to take. <laughter>

Coolures: Right.

Whitney: And none of us foresaw what digital images were going to do. None of us foresaw the Internet. And so the world has changed, but having been a part of it at this time was terrific.

Porter: Well, without your efforts, the things you just mentioned wouldn't have been possible.

Whitney: That's right.

Warner: And I think we're fortunate that IBM at that time was large enough to be able to develop technologies that spilled out into the industry. I mean, lots of people started at IBM and formed other businesses. But it allowed a growth, not only within the company, but around the world.

Lucke: Yeah, just look at the terminology. Winchester technology. I mean, we had a Winchester program, and now it's a technology that people have talked about for 20 years.

Porter: It's the one IBM code name that became part of a language.

All: <agreed>

END OF INTERVIEW

